

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 160 646 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
05.12.2001 Bulletin 2001/49

(51) Int Cl.7: **G06F 1/00**

(21) Application number: **01304493.8**

(22) Date of filing: **22.05.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: **Maria, Arturo**
Bellevue, WA 98004 (US)

(74) Representative: **Suckling, Andrew Michael**
Marks & Clerk
4220 Nash Court
Oxford Business Park South
Oxford OX4 2RU (GB)

(30) Priority: **30.05.2000 US 580689**

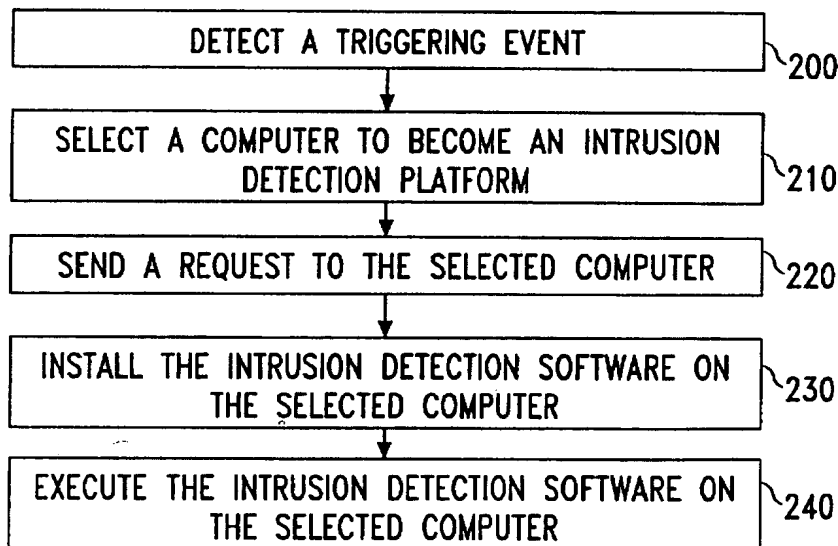
(71) Applicant: **AT&T WIRELESS SERVICES, INC.**
Redmond, WA 98052 (US)

(54) Floating intrusion detection platforms

(57) The present invention is a "floating" intrusion detection system that can use any computer on the network as an intrusion detection platform. A software agent program called a "socket" is installed on each computer that is to be available to be an intrusion detection platform. A central server contains intrusion detection software as well as a database containing knowledge based rules and profiles for detecting intrusions.

The central server can contact any computer that has a socket installed and direct that computer to become an intrusion detection platform. The selected computer then downloads, installs, and runs the intrusion detection software thus becoming an intrusion detection platform. Once the need has passed the central server can direct some of the platforms to stop running the software and return to their normal state.

FIG. 2



EP 1 160 646 A2

Description**FIELD OF THE INVENTION**

[0001] The present invention is directed to a method and system for providing dynamically distributed network security and intrusion detection.

BACKGROUND OF THE INVENTION

[0002] The importance of computer networks to companies' business interests and the interconnected nature of computer networks in the Internet era has resulted in increased concern about unauthorized network intrusions. When successful, these intrusions can cause damaging losses to the owner of the penetrated network in the form of vandalism, corporate espionage, theft of computer resources (when an intruder uses the penetrated network's computer resources for their own purposes, including attacking other networks), and negative publicity. Even just the potential of intrusion results in significant expenditures on computer resources to defend the network against intrusions including firewalls, proxy servers, and other intrusion detection and prevention systems.

[0003] Intrusion detection platforms are known. They are specialized hardware or software systems that use knowledge based rules and artificial intelligence concepts to detect attacks on computer networks so that defensive action can be taken. Examples of software used to implement intrusion detection platforms include Computer Associates' SessionWall, Check Point Software's RealSecure, and NetworkICE's BlackICE.

[0004] One type of intrusion detection system uses intrusion detection platforms placed at the entry points to networks where they inspect incoming network packets for signs that the packets are being employed in an attack on the network. If an attack is detected the intrusion detection platform may take several actions including alerting the system users, and refusing to allow the packets to enter the network. A primary drawback of these systems is that they require valuable computer hardware to be diverted from other uses and dedicated to simply monitoring and preventing intruders. Furthermore, in order to protect against insiders, such as disgruntled employees, these intrusion detection platforms generally must be distributed throughout the network in order to provide protection for the entire network, and in the event of a large scale attack or an attack localized to a particular area of the network, it is difficult to add new platforms or relocate existing platforms on short notice.

[0005] Another type of intrusion detection system resides on every computer in a network, and every computer monitors its own network security and reports back to a centralized server. These systems also have drawbacks because a portion of the processing power on every computer is dedicated to intrusion detection re-

sulting in a loss of performance to every user.

SUMMARY OF THE INVENTION

[0006] The present invention is a "floating" intrusion detection system that can dynamically change which computers on the network are acting as intrusion detection platforms. A software agent program called a "socket" is installed on each computer that is to be available to be an intrusion detection platform. A central server contains intrusion detection software as well as a database containing knowledge based rules and profiles for detecting intrusions. The central server can contact any computer that has a socket installed and direct that computer to become an intrusion detection platform. The selected computer then downloads, installs, and runs the intrusion detection software thus becoming an intrusion detection platform. The present invention allows the system to respond to network attacks or to simply respond to increases in network traffic by increasing the number of intrusion detection platforms whenever necessary. Once the need has passed, the central server can direct some of the platforms to stop running the software and return to their normal state. If a particular segment of the network is being attacked, more intrusion detection platforms could be added in that area without effecting other areas of the network. The present invention also allows a company to make more efficient use of their computer hardware. A computer that is used for a print server or scanner station during the work day could become an intrusion detection system at night without any human direction.

BRIEF DESCRIPTION OF THE DRAWINGS**[0007]**

Fig. 1 illustrates an example of a network in which the present invention might be implemented.

Fig. 2 is a flow chart illustrating one possible implementation of the method of the present invention.

Fig. 3 is a flow chart illustrating another possible implementation of the method of the present invention including a stop condition.

Fig. 4 is a flow chart illustrating one possible implementation of the method of the present invention.

Fig. 5 is a flow chart illustrating another possible implementation of the method of the present invention including a stop condition.

DETAILED DESCRIPTION

[0008] Fig. 1 is a diagram of an exemplary network suitable for use with the present invention. Network 103 may be any conventional network for data transmission including, for example, Ethernet, token ring, or RF hardware using TCP/IP, IPv6, or another appropriate network protocol. Network 103 may also include connec-

tions to other networks, including the Internet, via, for example, a direct connection (Hub 119) or a dial up connection (Modem 118) and typically employs a firewall 120 as a first line of defense against network intrusions. Connected to network 103 are servers 101 and 104 which may be conventional file servers capable of executing intrusion detection server software and may include databases 102 and 105. Connected to network 103 may also be a variety of typical computers (108 - 111, 114, and 115) and workstations (113, 116, and 117), some of which may also be connected to printers (112), scanners (118), or other peripheral equipment. These computers and workstations may also be separated into network segments 106 and 107. These network segments may be physically separated, logically separated, or both.

[0009] The deployment of the floating intrusion detection system of the present invention may be controlled or coordinated via a floating intrusion detection server (e.g., server 101). This server is equipped with a database that stores information about the network for which the server is detecting intrusions as well as a knowledgebase containing rules that define the server's operation including rules for identifying and responding to network intrusions, performing system maintenance, and scheduling predetermined system tasks. The information about the network that is stored in the database can include a network map, and/or a list of the computers within the network and their network addresses. Using this information the server can determine which computers in the network have been designated to be available for use as floating intrusion detection platforms. On each computer that has been so designated, there is a software agent program or "socket" running. The socket is a program that generally runs as a background process and listens for network messages from the floating intrusion detection server. The floating intrusion detection server can send messages to the socket at a computer instructing the socket to perform certain tasks including installing intrusion detection software, executing the intrusion detection software, and ceasing the execution of the intrusion detection software. The socket can also send messages back to the server containing information about the status of the computer.

[0010] As illustrated in Fig. 2, according to one embodiment of the present invention, when server 101 detects or is notified of a triggering event (Step 200), such as a possible network intrusion, the server selects an appropriate computer to become an intrusion detection platform (Step 210). The server then sends a request to the socket on that computer to become an intrusion detection platform (Step 220). The socket then installs (Step 230) and executes (Step 240) the intrusion detection software.

[0011] For example, server 101 may receive a message from firewall 120 indicating that an unusual number of incoming network packets directed at network segment 107 have been detected. In response to

this message, server 101, using the information about the network stored in database 102, selects computer 114, which is on segment 107, to become an intrusion detection platform. Server 101 then sends a message to the socket on computer 114, requesting that computer 114 become an intrusion detection platform. The socket on computer 114 receives the request, installs the intrusion detection software, and executes it. Thus an intrusion detection platform has been created that is at or near the target of the network attack.

[0012] Fig. 4 illustrates the actions taken by the socket on a remote computer according to one possible embodiment of the present invention. The socket receives a request from the intrusion detection server to become an intrusion detection platform (Step 400). The socket installs the intrusion detection software on the computer on which the agent is running (Step 410). The socket then executes the intrusion detection software and the computer begins functioning as an intrusion detection platform (Step 420).

[0013] The installation of the software on the computer may be accomplished in any number of ways. For example, the socket may download the software from a file server, the software may already be on the computer in a compressed archive, or the software may be attached to the request that came from the intrusion detection server. Additionally, the software installation may be accomplished in a multi-step process where components of the software are downloaded and installed from different locations. For example, the core software may be installed from a local archive and the latest update may be downloaded from a remote file server. Alternatively, the software may already be installed on the computer, and the socket only needs to check for software updates before executing the software.

[0014] The triggering event that causes the server to initiate new intrusion detection platforms may be defined by the administrator of the system, including, for example, increases or decreases in network traffic, unusual network traffic patterns, detection of network attacks by existing intrusion detection platforms or any other suspicious network activity. Additionally, the triggering event could simply be based on time of day, day of the week, etc. For example, since many network attacks occur after normal working hours, the system of the present invention could be configured to increase the number of intrusion detection platforms during these hours.

[0015] The intrusion detection server need not create more intrusion detection servers in response to every triggering event but may consider a number of factors before creating more platforms, including, for example, the number of intrusion detection platforms that already exist, the number of idle or underutilized eligible computers in the network, and predetermined minimum and maximum limits on the number of platforms.

[0016] As a complement to the triggering events that cause more intrusion detection platforms to be created,

the present invention also allows for "stop conditions" which are events or conditions that result in a computer ceasing execution of the intrusion detection software. These stop conditions may include, for example, the ceasing of the triggering event or condition that caused the intrusion detection platform to be created, a time period that has elapsed since the computer became an intrusion detection platform, or a request from a human operator. These "stop conditions" may be monitored or detected at the intrusion detection server which then sends a message to the intrusion detection platform instructing it to cease operating as an intrusion detection platform. Alternatively, the intrusion detection platform may monitor the stop condition itself and cease executing the intrusion detection software when the condition is fulfilled.

[0017] Triggering events and stop conditions may be specific to a particular computer or they may apply generally to all of the computers eligible to be intrusion detection platforms. For example, computer 111 may be designated to act as a print server for printer 112 during business hours and as an intrusion detection server after hours. Server 101 may have a triggering event and a stop condition specific to computer 111 in order to accomplish this schedule. Server 101 may also have a triggering event for a suspected network breach that directs server 101 to select any one (or more) of the eligible computers and request it to become an intrusion detection platform. Similarly, server 101 may maintain a stop condition that when a network attack ceases, server 101 selects a number of intrusion detection platforms and requests them to cease acting as intrusion detection platforms.

[0018] Fig. 3 illustrates one possible embodiment of the present invention for monitoring the stop condition at the intrusion detection server. Server 101 detects a triggering event (Step 300), selects an appropriate computer to become an intrusion detection platform, for example computer 110, (Step 310), and sends a request to the socket on computer 110 to become an intrusion detection platform (Step 320). Server 101 then monitors to detect if the stop condition has been fulfilled (Step 330). If the stop condition has not been fulfilled then server 101 continues to monitor, but if the stop condition has been fulfilled, server 101 sends a request to computer 110 to stop acting as an intrusion detection platform (Step 340).

[0019] Fig. 5 illustrates one possible embodiment of the present invention for monitoring the stop condition at the intrusion detection platform. The socket receives a request from the intrusion detection server to become an intrusion detection platform (Step 500). The socket executes the intrusion detection software and the computer begins functioning as an intrusion detection platform (Step 510). The socket and/or the intrusion detection software then monitors to see if the stop condition has been fulfilled (Step 520). This monitoring may be as simple as checking the date and time or the amount of

time the computer has been functioning as an intrusion detection platform, or may be more sophisticated monitoring of network traffic conditions. Once the stop condition has been fulfilled, the intrusion detection software ceases executing, the socket returns to the background and awaits further messages from the server (Step 530).

[0020] Some embodiments of the present invention may require that a number of messages be exchanged between the intrusion detection server and the sockets or intrusion detection software on the remote computers. In order to protect the intrusion detection system from being compromised by network attackers, these messages may be protected cryptographically. For example, the messages may be encrypted to prevent attackers from reading them, digitally signed to authenticate the sender, sent with a checksum or message digest to detect tampering, or any combination thereof. The encryption and digital signatures could use any of a number of well known techniques including RSA and DES. A number of secure checksum techniques are also known in the art.

[0021] To further protect the intrusion detection system from tampering or simple equipment failure, a secondary server may be employed in the system that maintains copies of the data on the primary server and immediately takes over if the primary server ceases operating correctly. This may be accomplished, for example, by server 101 sending updates to server 104 and database 105, or alternatively, server 104 could monitor server 101's network traffic in order to monitor server 101's activities.

[0022] The present invention is not limited to the specific embodiments described. It is expected that those skilled in the art will be able to devise other implementations that embody the principles of the present invention and remain within its scope.

Claims

1. A method for implementing an intrusion detection system in a network, comprising:

receiving a request at a software agent program to initiate intrusion detection services on a remote computer;
installing intrusion detection software on said remote computer via said software agent program; and
executing said intrusion detection software on said remote computer via said software agent program.

2. The method of claim 1 further comprising:
receiving a request to terminate intrusion detection services at said software agent program.

3. The method of claim 2 further comprising:

- monitoring for fulfillment of a stop condition.
4. The method of claim 3 wherein said stop condition is based on network traffic conditions. 5
 5. The method of claim 3 wherein said stop condition is an expiration time.
 6. The method of claim 1 further comprising the step of: 10
receiving notification of a network intrusion.
 7. The method of claim 6 further comprising the step of: 15
selecting said remote computer from a plurality of eligible computers.
 8. The method of claim 7 wherein said selecting step is accomplished based on a network map. 20
 9. The method of claim 7 wherein said selecting step is accomplished based on a knowledge base.
 10. The method of claim 1 wherein said request is verified using a cryptographic authentication scheme. 25
 11. The method of claim 1 wherein said request includes a stop condition indicating when to stop executing the intrusion detection software. 30
 12. The method of claim 11 wherein said stop condition is an expiration time.
 13. The method of claim 11 wherein said stop condition is based on network traffic conditions. 35
 14. The method of claim 7 wherein said request is verified using a cryptographic authentication scheme.
 15. A method for implementing an intrusion detection system on a computer connected to a network, comprising: 40
receiving a request to become an intrusion detection platform from a remote network location; and 45
executing said intrusion detection software.
 16. The method of claim 15 further comprising: 50
installing intrusion detection software on said computer.
 17. The method of claim 15 wherein said request includes a stop condition indicating when to stop executing the intrusion detection software. 55
 18. The method of claim 17 wherein said stop condition is an expiration time.
 19. The method of claim 17 wherein said stop condition is based on network traffic conditions.
 20. The method of claim 17 further comprising the step of: 60
when said stop condition is fulfilled, ceasing execution of said intrusion detection software.
 21. The method of claim 20 wherein said request is verified using a cryptographic authentication scheme.
 22. The method of claim 20 further comprising the step of: 65
when said intrusion detection software has ceased executing, un-installing said intrusion detection software.
 23. A system for detecting intrusions in a computer network comprising: 70
a plurality of computers executing software agents; 75
an intrusion detection server; and 80
a database, 85
wherein said intrusion detection server sends a request to execute intrusion detection software to a software agent at at least one of said plurality of computers when intrusion detection services are needed based on information contained in said database.
 24. The system of claim 23 wherein said intrusion detection server increases the number of said plurality of computers executing intrusion detection software when a network intrusion is detected.
 25. The system of claim 23 wherein said intrusion detection server changes the number of said plurality of computers executing intrusion detection software when the level of network traffic changes.
 26. The system of claim 23 wherein said intrusion detection server changes the number of said plurality of computers executing intrusion detection software depending on the time of day.
 27. The system of claim 23 wherein said database contains information about the plurality of computers.
 28. The system of claim 27 wherein said information includes a map of said computer network.
 29. The system of claim 23 wherein said database contains a knowledgebase.
 30. An article of manufacture comprising a computer-readable medium having stored thereon instructions adapted to be executed by a processor, the

Instructions which, when executed, define a series of steps to be used to perform network intrusion detection, said steps comprising:

receiving a request at a software agent program to initiate intrusion detection services on a remote computer; 5
installing intrusion detection software on said remote computer via said software agent program; and 10
executing said intrusion detection software on said remote computer.

31. The article of manufacture of claim 30 further comprising the step of: 15
receiving notification of a network intrusion.
32. The article of manufacture of claim 31 further comprising the step of: 20
selecting said remote computer from a plurality of eligible computers.
33. The article of manufacture of claim 32 wherein said selecting step is accomplished based on a network map. 25
34. The article of manufacture of claim 32 wherein said selecting step is accomplished based on a knowledge base. 30
35. The article of manufacture of claim 30 wherein said request is verified using a cryptographic authentication scheme.
36. The article of manufacture of claim 30 wherein said request includes a stop condition indicating when to stop executing the intrusion detection software. 35
37. The article of manufacture of claim 36 wherein said stop condition is an expiration time. 40
38. The article of manufacture of claim 36 wherein said stop condition is based on network traffic conditions. 45

50

55




FIG. 1

FIG. 2

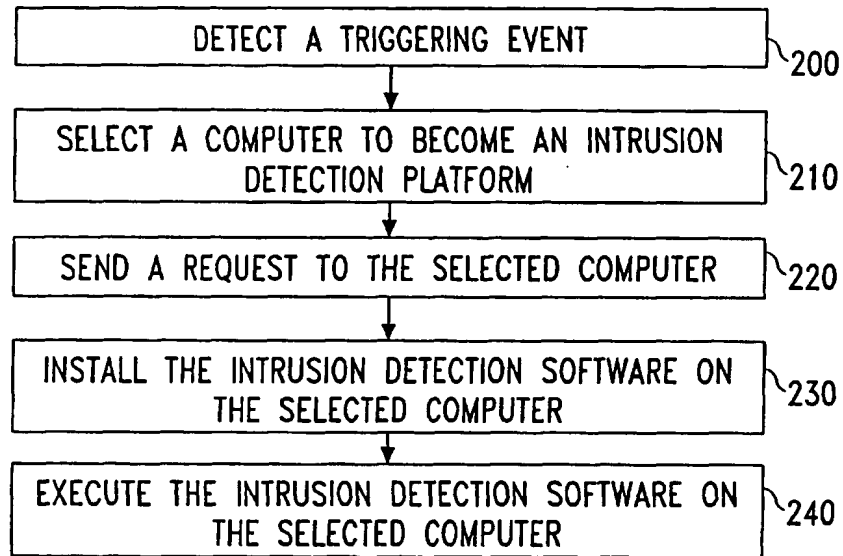


FIG. 3

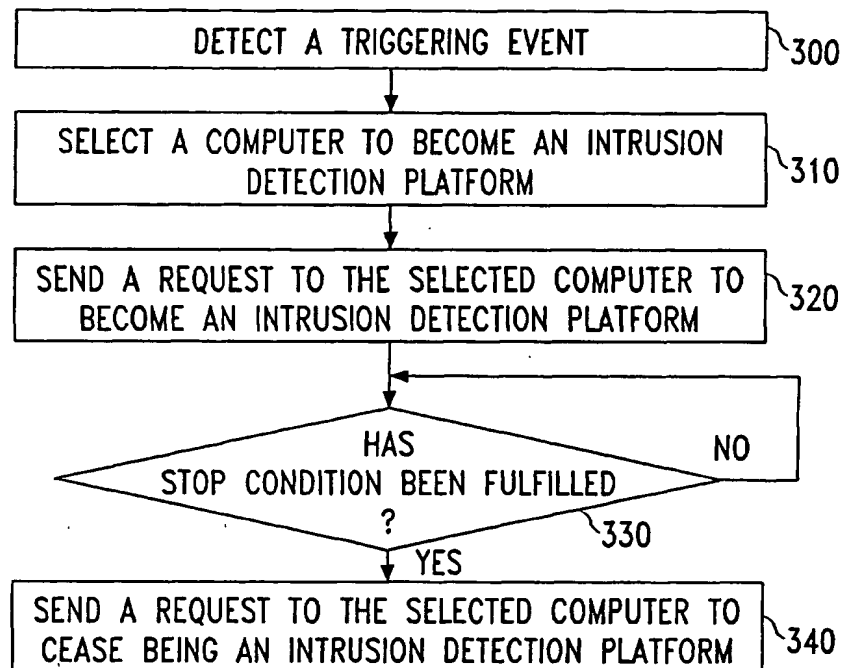


FIG. 4

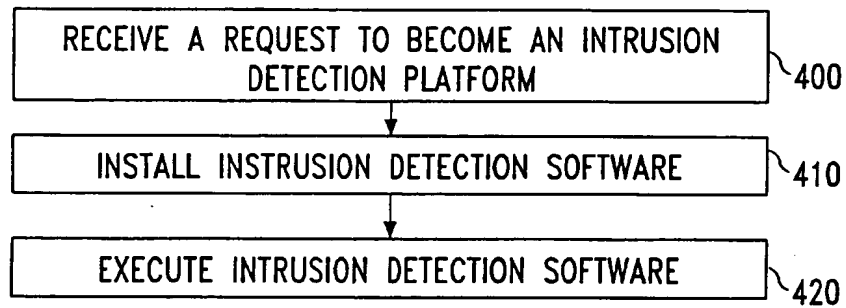


FIG. 5

